
The Max K -Armed Bandit: A PAC Lower Bound and tighter Algorithms

Yahel David and Nahum Shimkin
 Department of Electrical Engineering
 Technion - Israel Institute of Technology
 Haifa 32000, Israel

Abstract

We consider the Max K -Armed Bandit problem, where a learning agent is faced with several sources (arms) of items (rewards), and interested in finding the best item overall. At each time step the agent chooses an arm, and obtains a random real valued reward. The rewards of each arm are assumed to be i.i.d., with an unknown probability distribution that generally differs among the arms. Under the PAC framework, we provide lower bounds on the sample complexity of any (ϵ, δ) -correct algorithm, and propose algorithms that attain this bound up to logarithmic factors. We compare the performance of this multi-arm algorithms to the variant in which the arms are not distinguishable by the agent and are chosen randomly at each stage. Interestingly, when the maximal rewards of the arms happen to be similar, the latter approach may provide better performance.

1 Introduction

In the classic stochastic multi-armed bandit (MAB) problem the learning agent faces a set K of stochastic arms, and wishes to maximize its cumulative reward (in the regret formulation), or find the arm with the highest mean reward (the pure exploration problem). This model has been studied extensively in the statistical and learning literature, see for example [1] for a comprehensive survey.

We consider a variant of the MAB problem called the Max K -Armed Bandit problem (Max-Bandit for short). In this variant, the objective is to obtain a sample with the highest possible reward (namely, the highest value in the support of the probability distribution of any arm). More precisely, considering the PAC setting, the objective is to return an (ϵ, δ) -correct sample, namely a sample which its reward value is ϵ -close to the overall best possible reward with a probability larger than $1 - \delta$. In addition, we wish to minimize the sample complexity, namely the expected number of samples observed by the learning algorithm before it terminates.

For the classical MAB problem, algorithms that find the best arm (in terms of its expected reward) in the PAC sense were presented in [2, 3, 4], and lower bounds on the sample complexity were presented in [5] and [3]. The essential difference with respect to this work is in the objective, which is to find an (ϵ, δ) -correct sample in our case. The scenario considered in Max-Bandit model is relevant when a single best item needs to be selected from among several (large) clustered sets of items, with each set represented as a single arm. These sets may represent parts that come from different manufacturers or produced by different processes, job candidates that are referred by different employment agencies, finding the best match to certain genetic characteristics in different populations, or choosing the best channel among different frequency bands in a cognitive radio wireless network.

The Max-Bandit problem was apparently first proposed in [6]. For reward distribution functions in a specific family, an algorithm with an upper bound on the sample complexity that increases as $\frac{-\ln(\delta)}{\epsilon^2}$ was provided in [7]. For the case of discrete rewards, another algorithm was presented in

[8], without performance analysis. Later, a similar model in which the objective is to maximize the expected value of the largest sampled reward for a given number of samples (n) was studied in [9]. In that work the attained best reward is compared with the expected reward obtained by an oracle that samples the best arm n time. An algorithm is suggested and shown to secure an upper bound of order $n^{-\alpha}$ on that difference, where $\alpha < 1$ is determined by the properties of the distribution functions and decreases as they are further away from a specific functions family.

Our basic assumption in the present paper is that a known lower bound is available on the tail distributions, namely on the probability that the reward of each given arm will be close to its maximum. A special case is when the probability densities near the maximum are larger than a given value, but we consider more general function classes. Under that assumption, we provide an algorithm for which the sample complexity increases as at most $\frac{-\ln(\delta) - \ln(\epsilon)}{\epsilon}$. This provides an improvement by a factor of ϵ^{-1} over the result of [7], which was obtained for a more specific model. To compare with the result in [9], we observe that with a choice of $\delta = \frac{1}{n^2}$ in our algorithm, we obtain that the expected shortfall of the largest sample with respect to the maximal reward possible is at most of order $O(\frac{\ln(n)}{n})$ (as compared to $O(n^{-\alpha})$ with $\alpha < 1$). Furthermore, we provide a lower bound on the sample complexity of every (ϵ, δ) -correct algorithm, which holds when several arms posses maximal rewards that are close to that of the best arm. This lower bound is shown to coincide, up to a logarithmic term, with the upper bound derived for the proposed algorithm.

A basic feature of the Max-Bandit problem (and the associated algorithms) is the goal of quickly focusing on the best arm (in term of maximal reward), and sampling from that arm as much as possible. It should be of interest to compare the obtained results with the alternative approach, which ignores the distinction between arms, and simply draws a sample from a random arm (say, with uniform probabilities) at each round. This can be interpreted as mixing the items associated with each arm before sampling; we accordingly refer to this variant as the unified-arm problem. This problem actually coincides with the so-called infinitely-many armed bandit model studied in [10, 11, 12, 13, 14], for the specific case of deterministic arms studied in [15]. The conclusion about weather to apply the multi-arm approach or the unified-arm approach is inconclusive. However, as a rule of thumb, when the maximal possible rewards of many arms are far from the optimal, the multi-arm approach has better performance.

The paper proceeds as follows. In the next section we present our model. In Section 3 we provide a lower bound on the sample complexity of every (ϵ, δ) -correct algorithm. In Section 4 we present two (ϵ, δ) -correct algorithms, and we provide an upper bound on the sample complexity of one of them. The first algorithm is simple and its bound has the same order as the lower bound up to a logarithmic term in $\frac{|K|}{\epsilon}$ (where $|K|$ stands for the number of arms), the second algorithm is more complicated and we believe that its bound is larger by up to a double logarithmic term in $\frac{|K|}{\epsilon}$ than the lower bound. In Section 5, we consider for comparison the unified-arm case. In Section 6 we close the paper by some concluding remarks. Certain proofs are deferred to the Appendix due to space limitations.

2 Model Definition

We consider a finite set of arms, denoted by K . At each stage $t = 1, 2, \dots$ the learning agent chooses an arm $k \in K$, and a real valued reward is obtained from that arm. The rewards obtained from each arm k are independent and identically distributed, with a distribution function (CDF) $F_k(\mu)$, $\mu \in \mathbb{R}$. We denote the maximal possible reward of each arm by $\mu_k^* = \inf_{\mu \in \mathbb{R}} \{\mu | F_k(\mu) = 1\}$, assumed finite, and the maximal reward among all arms by $\mu^* = \max_{k \in K} \mu_k^*$.

Throughout the paper, we shall make the following assumption.

Assumption 1. *There exist known constants $A > 0$, $\beta \geq 0$ and $\epsilon_0 > 0$ such that, for every $k \in K$ and $0 < \epsilon \leq \epsilon_0$, it holds that*

$$P(\mu_k > \mu_k^* - \epsilon) \geq A\epsilon^\beta,$$

where μ_k stands for a random variable with distribution F_k .

The bound in the above assumption can also be expressed as $1 - F_k(\mu) \geq A(\mu_k^* - \mu)^\beta$. This condition required μ_k to have a certain mass near its maximal reward. Note that the specific case

of $\beta = 1$ is satisfied if the densities F'_k are lower bounded by a constant A . Values of $\beta < 1$ accommodate leaner tales.

The upper bound on the CDF F_k ensures that for each arm, an ϵ -optimal reward can be observed by a finite number of samples. The bound in the above assumption is similar to those assumed in [12] and [15].

An algorithm for the Max-Bandit model samples an arm at each time step, based on the observed history so far (i.e., the previously selected arms and observed rewards). We require the algorithm to terminate after a random number T of samples, which is finite with probability 1, and return a reward V which is the maximal reward observed over the entire period. An algorithm is said to be (ϵ, δ) -correct if

$$P(V > \mu^* - \epsilon) > 1 - \delta.$$

The expected number of samples $E[T]$ taken by the algorithm is the *sample complexity*, which we wish to minimize.

3 A Lower Bound

Before turning to our proposed algorithm, we provide a lower bound on the sample complexity of any (ϵ, δ) -correct algorithm. The bound holds under Assumption 1 when $\beta \leq 1$. The case of $\beta > 1$ is more complicated for analysis and it is still unclear whether our lower bound holds for this case.

The following result specifies the lower bound of this section.

Theorem 1. *Suppose $\epsilon_0 \leq (4A)^{-1/\beta}$, $\beta \leq 1$ and let $\epsilon \in (0, \epsilon_0)$ and $\delta \in (0, 1)$. Let k^* denote some optimal arm, such that $\mu_{k^*}^* = \mu^*$. Then, under Assumption 1, for every (ϵ, δ) -correct algorithm, it holds that*

$$E[T] \geq \sum_{k \in K \setminus \{k^*\}} \frac{1}{8A(\min(\epsilon_0, \epsilon + \mu^* - \mu_k^*))^\beta} \ln\left(\frac{3}{16\delta}\right). \quad (1)$$

This lower bound can be interpreted as summing over the minimal number of times that each arm, other than the optimal arm k^* , needs to be sampled. It is important to observe that if there are several optimal arms, only one of them is excluded from the summation. Indeed, the bound is most effective when there are several optimal (or near-optimal) arms, as the denominator of the summand is larger for such arms. This may appear surprising at first, as more sources of good rewards are available; however, when there is a single arm that is strictly better than the others it can be quickly singled out, while if many arms have nearly optimal rewards, more samples are "wasted" on determining which arm is best.

The proof of Theorem 1 is provided in Appendix A and proceeds by showing that if an algorithm is (ϵ, δ) -correct and its sample complexity is lower than a certain threshold for some set of reward distributions, then this algorithm cannot be (ϵ, δ) -correct for some related reward distributions.

4 Algorithms

Here we provide two (ϵ, δ) -correct algorithms. The first algorithm is based on sampling the arm which has the highest upper confidence bound on its maximal reward at each time step and the second algorithm is based on arms elimination.

4.1 Maximal Confidence Bound

The algorithm starts by sampling a certain number of times from each arm. Then, it repeatedly calculates an index for each arm which can be interpreted as a certain upper bound on the maximal reward of this arm, and samples once from the arm with the largest index. The algorithm terminates when the number of samples from the arm with the largest index is above a certain threshold. This idea is similar to that in the UCB1 Algorithm provided in [16].

Algorithm 1 Maximal Confidence Bound (Max-CB) Algorithm

- 1: **Input:** Model parameters $\epsilon_0 > 0$, $A > 0$ and $\beta \geq 0$, constants $\delta > 0$ and $\epsilon > 0$.
Define $L = 6 \ln \left(|K| \left(1 + \frac{-\ln(\delta)}{A\epsilon^\beta} \right) \right)$.
- 2: **Initialization:** Counters $C(k) = N_0$, $k \in K$, where $N_0 = \lfloor \frac{L - \ln(\delta)}{A\epsilon_0^\beta} \rfloor + 1$.
- 3: Sample N_0 times from each arm.
- 4: Compute $Y_{C(k)}^k = V_{C(k)}^k + \epsilon^{UB}(C(k))$ and set $k^* \in \arg \max_{k \in K} Y_{C(k)}^k$ (with tie broken arbitrary), where $V_{C(k)}^k$ is the largest reward observed so far from arm k and

$$\epsilon^{UB}(C(k)) = \left(\frac{L - \ln(\delta)}{AC(k)} \right)^{1/\beta}.$$

- 5: If $\epsilon^{UB}(C(k^*)) < \epsilon$, stop and return the largest sampled reward.
Else, sample once from arm k^* , set $C(k^*) = C(k^*) + 1$ and return to step 4.
-

Theorem 2. Under Assumption 1, for $L \geq 10$, Algorithm 1 is (ϵ, δ) -correct with a sample complexity of

$$E[T] \leq \sum_{k \in K} \frac{L - \ln(\delta)}{A(\max(\epsilon, \mu^* - \mu_k^*))^\beta} + |K|N_0,$$

where $N_0 = \lfloor \frac{L - \ln(\delta)}{A\epsilon_0^\beta} \rfloor + 1$ and $L = 6 \ln \left(|K| \left(1 + \frac{-\ln(\delta)}{A\epsilon^\beta} \right) \right)$ as defined in the algorithm.

In the following corollary we present the ratio between the lower bound presented in Theorem 1 to the upper bound in Theorem 2.

Corollary 1. If there are more than one arm for which $\mu_k^* \in [\mu^* - \epsilon, \mu^*]$, then the upper bound on the sample complexity is of the same order as the lower bound in Theorem 1, up to a logarithmic factor in $\frac{|K|}{\epsilon}$.

Proof. For every $k \in K$ it follows that

$$\Theta_k^1 \triangleq \frac{1 + 2^\beta}{(\min(\epsilon_0, \epsilon + \mu^* - \mu_k^*))^\beta} \geq \frac{2^\beta}{(\epsilon + \mu^* - \mu_k^*)^\beta} + \frac{1}{\epsilon_0^\beta} \geq \frac{1}{(\max(\epsilon, \mu^* - \mu_k^*))^\beta} + \frac{1}{\epsilon_0^\beta} \triangleq \Theta_k^2,$$

and for every two arms k' and k^* for which $\mu_{k'}^* \in [\mu^* - \epsilon, \mu^*]$ and $\mu_{k^*}^* = \mu^*$ it is obtained that

$$\Theta_{k'}^1 \geq 2^{-\beta} \Theta_{k^*}^1. \quad (2)$$

In addition, the lower bound is of the same order as

$$-\ln(\delta) \sum_{k \in K \setminus \{k^*\}} \Theta_k^1, \quad (3)$$

the upper bound is of the same order as

$$(L - \ln(\delta)) \sum_{k \in K} \Theta_k^2,$$

Therefore, the upper bound in Theorem 2 is of the same order of the lower bound in Theorem 1 up to an order of $\frac{L - \ln(\delta)}{-\ln(\delta)}$, which is logarithmic in $\frac{|K|}{\epsilon}$. □

To establish Theorem 2, we first bound the probability of the event under which the upper bound of the best arm is below the maximal reward. Then, we bound the largest number of samples after which the algorithm terminates under the assumption that the upper bound of the best arm is above the maximal reward.

Proof (Theorem 2). We denote the time step of the algorithm by t , and the value of the counter $C(k)$ at time step t by $C^t(k)$. Recall that T stands for the random final time step. By the condition in step 5 of the algorithm, for every arm $k \in K$, it follows that,

$$C^T(k) \leq \lfloor \frac{L - \ln(\delta)}{A\epsilon^\beta} \rfloor + 1. \quad (4)$$

Note that by the fact that for $x \geq 6$ it follows that $\frac{d6 \ln(x)}{dx} \leq 1$, and by the fact that for $x_0 = \exp(1\frac{2}{3})$ it follows that $x_0 > 6 \ln(x_0) = 10$ it is obtained that

$$L' \triangleq |K| \left(\frac{-\ln(\delta)}{A\epsilon^\beta} + 1 \right) > 6 \ln \left(|K| \left(\frac{-\ln(\delta)}{A\epsilon^\beta} + 1 \right) \right) = L,$$

for $L \geq 10$. So, by the fact that $T = \sum_{k \in K} C^T(i)$, for $L \geq 10$ it follows that

$$T \leq |K| \left(\frac{L - \ln(\delta)}{A\epsilon^\beta} + 1 \right) < |K| \left(\frac{L' - \ln(\delta)}{A\epsilon^\beta} + 1 \right) \leq L'^2 = e^{\frac{L}{3}}. \quad (5)$$

Now, we begin with proving the (ϵ, δ) -correctness property of the algorithm. Recall that for every arm $k \in K$ the rewards are distributed according to the C.D.F. $F_k(\mu)$. Let assume w.l.o.g. that $\mu_1^* = \mu^*$. Then, for $N > 0$ and by the fact that $(1 - \epsilon)^{\frac{1}{\epsilon}} \leq e^{-1}$ for every $\epsilon \in (0, 1]$, for $\epsilon^{UB}(N) = \left(\frac{L - \ln(\delta)}{AN} \right)^{1/\beta}$ it follows that

$$P(V_N^1 \leq \mu^* - \epsilon^{UB}(N)) = (F_1(\mu^* - \epsilon^{UB}(N)))^N \leq (1 - A(\epsilon^{UB}(N))^\beta)^N \leq \delta e^{-L}, \quad (6)$$

where V_N^k is the largest reward observed from arm $k \in K$ after this arm has been sampled for N times. Hence, at every time step t , by the definition of $Y_{C^t(1)}^1$ and Equations (5) and (6), by applying the union bound, it follows that

$$P(Y_{C^t(1)}^1 \leq \mu^*) \leq P(V_{C^t(1)}^1 \leq \mu^* - \epsilon^{UB}(C^t(1))) \leq \sum_{t=1}^{\exp(\frac{L}{3})} P(V_N^1 \leq \mu^* - \epsilon^{UB}(N)) \leq \delta e^{-\frac{2L}{3}}. \quad (7)$$

Since by the condition in step 5, it is obtained that when the algorithm stops

$$V_{C^t(k^*)}^{k^*} > Y_{C^t(k^*)}^{k^*} - \epsilon,$$

and by the fact that for every time step

$$Y_{C^t(k^*)}^{k^*} \geq Y_{C^t(1)}^1,$$

it follows by Equation (7) that

$$P(V_{C^t(k^*)}^{k^*} \leq \mu^* - \epsilon) \leq P(Y_{C^t(1)}^1 \leq \mu^*) \leq \delta e^{-\frac{2L}{3}}.$$

Therefore, it follows that the algorithm returns a reward greater than $\mu^* - \epsilon$ with a probability larger than $1 - \delta$. So, it is (ϵ, δ) -correct.

For proving the bound on the expected sample complexity of the algorithm we define the following sets:

$$M(\epsilon) = \{l \in K | \mu_l^* - \mu_l^* < \epsilon\}, \quad N(\epsilon) = \{l \in K | \mu_l^* - \mu_l^* \geq \epsilon\}.$$

As before, we assume w.l.o.g. that $\mu_1^* = \mu^*$. For the case in which

$$E_1 \triangleq \bigcap_{1 \leq t < T} \{Y_{C^t(1)}^1 \geq \mu^*\},$$

occurs, since $V_{C^t(k)}^k \leq \mu_k^*$ for every $k \in K$, and every time step, it follows that the necessary condition for sampling from arm k ,

$$Y_{C^k(1)}^k \geq Y_{C^t(1)}^1,$$

occurs only when the event

$$E_2(t) \triangleq \{ \mu_k^* + \epsilon^{UB} (C^t(k)) \geq \mu^* \},$$

occurs. But

$$E_2(t) \subseteq \left\{ C^t(k) \leq \frac{L - \ln(\delta)}{A(\mu^* - \mu_k^*)^\beta} \right\}.$$

Therefore, it is obtained that

$$C^T(k) \leq \lfloor \frac{L - \ln(\delta)}{A(\mu^* - \mu_k^*)^\beta} \rfloor + 1. \quad (8)$$

By using the bound in Equation (4) for the arms in the set $M(\epsilon)$, the bound in Equation (8) for the arms in the set $N(\epsilon)$ and the bound in Equation (5), it is obtained that

$$E[T] \leq (1 - P(E_1)) e^{\frac{L}{3}} + P(E_1) \Phi(\epsilon), \quad (9)$$

where

$$\Phi(\epsilon) \triangleq \left(\sum_{k \in N(\epsilon)} \left(\lfloor \frac{L - \ln(\delta)}{A(\mu^* - \mu_k^*)^\beta} \rfloor + 1 \right) + \sum_{k \in M(\epsilon)} \left(\lfloor \frac{L - \ln(\delta)}{A\epsilon^\beta} \rfloor + 1 \right) \right).$$

In addition, by Equation (7), the bound in Equation (5) and by applying the union bound, it follows that

$$P(E_1) \geq 1 - \sum_{t=1}^T P(Y_{C^t(1)}^1 < \mu^*) \geq 1 - \delta e^{-\frac{2L}{3}} e^{\frac{L}{3}} = 1 - \delta e^{-\frac{L}{3}}.$$

So,

$$1 - P(E_1) \leq \delta e^{-\frac{L}{3}}. \quad (10)$$

Furthermore, by the definitions of the sets $N(\epsilon)$ and $M(\epsilon)$, it can be obtained that

$$\Phi(\epsilon) \leq \sum_{k \in K} \lfloor \frac{L - \ln(\delta)}{A(\max(\epsilon, \mu^* - \mu_k^*))^\beta} \rfloor + 1. \quad (11)$$

Therefore, by Equation (9), (10) and (11) the bound on the sample complexity is obtained. \square

4.2 Maximal Eliminator

The algorithm starts by sampling a certain number of times from each arm. Then, it repeatedly calculates an index for each arm which can be interpreted as a certain upper bound on the maximal reward of this arm, and eliminates arms for which that index is below the maximal sampled reward so far. Then it sample from only the retained arms (those arms which have not been eliminated) a number of times that is doubled at each sampling phase. This idea is similar to that in the Median Elimination Algorithm provided in [2].

We do not provide performance analysis for Algorithm 2. However, since the number of times at which the confidence bounds should be correct (times at which the algorithm eliminates arms) is only logarithmic in the number of total samples, we have $L^{ME} = \ln(2L)$ (where L is defined in Algorithm 1 and the factor 2 arises because of the doubling). Therefore, we believe that the upper bound on the sample complexity of Algorithm 2 would be that of Algorithm 1 multiplied by $\frac{2L^{ME}}{L}$. So, the upper bound would be of the same order of the lower bound in Theorem 1 up to double logarithmic terms.

5 Comparison with The Unified-Arm Model

In this section, we analyze the improvement in the sample complexity obtained by utilizing the multi arm property (the ability to choose from which arm to sample at each time step) compared

Algorithm 2 Maximal Eliminator (ME) Algorithm

- 1: **Input:** Model parameters $\epsilon_0 > 0$, $A > 0$ and $\beta \geq 0$, constants $\delta > 0$ and $\epsilon > 0$.
Define $L^{ME} = \ln \left(12 \ln \left(|K| \left(1 + \frac{-\ln(\delta)}{A\epsilon^\beta} \right) \right) \right)$.
 - 2: **Initialization:** A set of arms $K_{t=1} = K$ and counter $t = 1$.
 - 3: Sample N_t times from each arm in the set K_t , where $N_t = 2^{t-1} \left(\lfloor \frac{L^{ME} - \ln(\delta)}{A\epsilon_0^\beta} \rfloor + 1 \right)$.
 - 4: Compute $Y^k = V^k + \epsilon^{UB}(N_{t+1} - N_0)$,
where V^k is the largest reward observed from arm k and $\epsilon^{UB}(N) = \left(\frac{L^{ME} - \ln(\delta)}{AN} \right)^{1/\beta}$.
 - 5: If $\epsilon^{UB}(N_{t+1} - N_0) < \epsilon$, stop and return the largest sampled reward.
Else, set $t = t + 1$, $K_t = \{k \in K_{t-1} | Y^k \geq \max_{j \in K_{t-1}} V^j\}$ and return to step 3.
-

Algorithm 3 Unified-Arm Algorithm

- 1: **Input:** Constants $\delta > 0$, $\epsilon > 0$.
 - 2: Sample $\lceil \frac{-\ln(\delta)|K|}{A\epsilon^\beta} \rceil + 1$ arms from the arm.
 - 3: Return the best sampled arm.
-

to a model in which all the arms are unified into a unified arm, so that the sample is effectively obtained from a random arm. In the unified-arm model, when the agent samples from this unified arm, a certain arm (among the multi arm) is chosen uniformly and a reward is sampled from this arm. We denote the CDF of the unified arm as $F(\mu)$, with $F = \frac{1}{|K|} \sum_{k \in K} F_k$. By Assumption 1, $1 - F(\mu) \geq \frac{A(\mu^* - \mu)^\beta}{|K|}$, and the corresponding maximal reward is μ^* .

In the remainder of this section, we provide a lower bound on the sample complexity and an (ϵ, δ) -correct algorithm that attains the same order of this bound for the unified-arm model. (Note that the lower bound in Theorem 1 is meaningless for $|K| = 1$.) Then, we discuss which approach (multi-arm or unified-arm) is better for different model parameters, and provide examples that illustrate these cases.

5.1 Lower Bound

The following Theorem provides a lower bound on the sample complexity for the unified-arm model.

Theorem 3. Suppose $\epsilon_0 \leq \left(\frac{|K|}{2A} \right)^{\frac{1}{\beta}}$, $\beta \leq 1$ and let $\epsilon \in (0, \epsilon_0)$, $\delta \in (0, 1)$. Then, under Assumption 1, for every (ϵ, δ) -correct algorithm, it holds that

$$E[T] \geq \frac{|K|}{4A\epsilon^\beta} \ln \left(\frac{3}{5\delta} \right). \quad (12)$$

The proof is provided in Appendix B and is based on the a similar idea to that of Theorem 1.

5.2 Algorithm

In Algorithm 3 a certain number of rewards is sampled, and the algorithm chooses the best one among them. In the following Theorem we provide a bound on the sample complexity achieved by Algorithm 3.

Theorem 4. Under Assumption 1, Algorithm 3 is (ϵ, δ) -correct, with a sample complexity bound of

$$E[T] \leq \frac{|K| \ln(\delta^{-1})}{A\epsilon^\beta} + 2.$$

The proof is provided in Appendix C. Note that the upper bound on the sample complexity is of the same order as the lower bound in Theorem 3.

5.3 Comparison and Examples

To find when the multi-arm algorithm is helpful, we can compare the upper bound on the sample complexity provided in Theorem 2 for Algorithm 1 (multi-arm case) with the lower bound for the unified-arm model in Theorem 3.

Case 1: Suppose first that arm 1 is best: $\mu_1^* = \mu^*$, while all the other arms fall short significantly compared to the required accuracy ϵ : $\mu_k^* \ll \mu^* - \epsilon$, for $k \neq 1$.

In this case $\frac{1}{\epsilon^\beta} \gg \frac{1}{(\max(\epsilon, \mu^* - \mu_k^*))^\beta}$, for $k \neq 1$. Hence the upper bound on sample complexity of Algorithm 1 (multi-arm case) will be smaller than the lower bound for the unified-arm model in Theorem 3. We now provide an example which illustrate case 1 numerically.

Example 1 (Case 1). *Let $|K| = 10^4$, $\mu_1^* = 0.9$, $\mu_k^* = 0.1 \forall k \in K \setminus \{1\}$, $\beta = 1$ and $A = 0.01$. For $\epsilon = 10^{-4}$ and $\delta = 10^{-3}$ the sample complexity attained by Algorithm 1 is 3.52×10^8 . The lower bound for the unified-arm model is 1.59×10^{10} . The sample complexity attained by Algorithm 3 (for the unified-arm model) is 6.9×10^{10} .*

Case 2: Consider next the opposite case, where there are many optimal arms and few that are worse: say $\mu_1^* \ll \mu^* - \epsilon$, while $\mu_k^* = \mu^*$ for all $k \neq 1$.

In this case $\frac{1}{\epsilon^\beta} = \frac{1}{(\max(\epsilon, \mu^* - \mu_k^*))^\beta}$, for $k \neq 1$. Hence, since there is a logarithmic-in- $\frac{|K|}{\epsilon}$ multiplicative factor in the upper bound on the sample complexity of Algorithm 1 (multi-arm case), this bound will be larger than the lower bound for the unified-arm model in Theorem 3. The following example illustrate case 2 numerically.

Example 2 (Case 2). *Let $|K|$, A , β , δ and ϵ remain the same as in Example 1, and let $\mu_1^* = 0.1$ and $\mu_k^* = 0.9 \forall k \in K \setminus \{1\}$. The sample complexity of Algorithm 1 is 1.56×10^{12} , which is larger than the sample complexity of Algorithm 3 which is 6.9×10^{10} .*

As shown in Example 2, in some cases the bound on the sample complexity of Algorithm 1 (multi-arm) is larger than that of Algorithm 3 (unified-arm). By comparing the upper bounds of these algorithms, we believe that the logarithmic in $\frac{|K|}{\epsilon}$ factor in the bound of Algorithm 1 may not be required.

As observed by comparing the lower and upper bounds for the multi-arm and the unified-arm model, the unified-arm algorithm provides a tighter upper bound (compared to the matching lower bound). Therefore, when the benefit obtained by the multi-arm model is small (i.e., when there are a lot of good arms) the profit obtained by applying the multi-arm Algorithm turns out to be loss.

6 Conclusion

In this paper we have developed corresponding lower and upper bounds on the sample complexity, which are essentially the same order up to a logarithmic term in $\frac{|K|}{\epsilon}$ for the Max K -Armed Bandit problem.

These results were compared to the unified-arm model, where the learning algorithm effectively unifies the different arms into one. While the multi-arm algorithm usually performs better, in some cases, in particular when most arms are optimal, the unified arm algorithm may provide better performance. It still remains to be shown whether an algorithm that provides the performance benefits of both approaches may be devised.

Another direction for future work concerns the relaxation or generalization of our Assumption 1, which requires a known lower bound on the tail distribution of the rewards.

References

- [1] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.
- [2] E. Even-Dar, S. Mannor, and Y. Mansour, "PAC bounds for multi-armed bandit and markov decision processes," in *Computational Learning Theory*, pp. 255–270, 2002.

- [3] J.-Y. Audibert and S. Bubeck, “Best arm identification in multi-armed bandits,” in *COLT-23th Conference on Learning Theory-2010*, pp. 13–p, 2010.
- [4] V. Gabillon, M. Ghavamzadeh, and A. Lazaric, “Best arm identification: A unified approach to fixed budget and fixed confidence,” in *Advances in Neural Information Processing Systems 25*, pp. 3212–3220, Curran Associates, Inc., 2012.
- [5] S. Mannor and J. N. Tsitsiklis, “The sample complexity of exploration in the multi-armed bandit problem,” *Journal of Machine Learning Research*, vol. 5, pp. 623–648, 2004.
- [6] V. A. Ciciirello and S. F. Smith, “The max k-armed bandit: A new model of exploration applied to search heuristic selection,” in *Proceedings of the National Conference on Artificial Intelligence*, vol. 20, p. 1355, 2005.
- [7] M. J. Streeter and S. F. Smith, “An asymptotically optimal algorithm for the max k-armed bandit problem,” in *Proceedings of the National Conference on Artificial Intelligence*, vol. 21, p. 135, 2006.
- [8] M. J. Streeter and S. F. Smith, “A simple distribution-free approach to the max k-armed bandit problem,” in *Principles and Practice of Constraint Programming-CP 2006*, pp. 560–574, Springer, 2006.
- [9] A. Carpentier and M. Valko, “Extreme bandits,” in *Advances in Neural Information Processing Systems 27*, pp. 1089–1097, Curran Associates, Inc., 2014.
- [10] D. A. Berry, R. W. Chen, A. Zame, D. C. Heath, and L. A. Shepp, “Bandit problems with infinitely many arms,” *The Annals of Statistics*, pp. 2103–2116, 1997.
- [11] O. Teytaud, S. Gelly, and M. Sebag, “Anytime many-armed bandits,” in *CAP*, (Grenoble, France), 2007.
- [12] Y. Wang, J.-Y. Audibert, and R. Munos, “Infinitely many-armed bandits,” *Advances in Neural Information Processing Systems*, vol. 8, pp. 1–8, 2008.
- [13] D. Chakrabarti, R. Kumar, F. Radlinski, and E. Upfal, “Mortal multi-armed bandits,” in *Advances in Neural Information Processing Systems 21*, pp. 273–280, Curran Associates, Inc., 2009.
- [14] T. Bonald and A. Proutiere, “Two-target algorithms for infinite-armed bandits with Bernoulli rewards,” in *Advances in Neural Information Processing Systems 26*, pp. 2184–2192, Curran Associates, Inc., 2013.
- [15] Y. David and N. Shimkin, “Infinitely many-armed bandits with unknown value distribution,” in *Machine Learning and Knowledge Discovery in Databases*, pp. 307–322, Springer, 2014.
- [16] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Machine Learning*, vol. 47, pp. 235–256, 2002.

7 Appendix A

Proof (Theorem 1). Let $\bar{\mu}_k = \sup_{\mu \in \mathbb{R}} \{\mu | F_k(\mu) \leq 1 - A\epsilon_0^\beta\}$ for every $k \in K$. Then, we define the following set of hypotheses $\{H_0, H_1, \dots, H_{|K|}\}$:

$$H_0 : f_k^{H_0}(\mu) = f_k(\mu) \quad \forall k \in K,$$

and, for every $k = 1, \dots, |K|$,

$$\begin{aligned} H_k : f_l^{H_k}(\mu) &= f_l(\mu), \quad l \neq k, \\ \text{if } \epsilon_0 \leq (\mu^* - \mu_k^* + \epsilon) : f_k^{H_k}(\mu) &= \gamma_k^1 f_k(\mu) \mathbf{1}_{(-\infty, \mu_k^*)}(\mu) + A\beta (\mu^* + \epsilon - \mu)^{\beta-1} \mathbf{1}_{(\mu^* + \epsilon - \epsilon_0, \mu^* + \epsilon]}(\mu), \\ \text{if } \epsilon_0 > (\mu^* - \mu_k^* + \epsilon) : f_k^{H_k}(\mu) &= \gamma_k^2 f_k(\mu) \mathbf{1}_{(-\infty, \bar{\mu}_k)}(\mu) + \gamma_k^3 f_k(\mu) \mathbf{1}_{(\mu = \bar{\mu}_k)} + f_k(\mu) \mathbf{1}_{(\bar{\mu}_k, \mu_k^*]}(\mu) \\ &\quad + A\beta (\mu^* + \epsilon - \mu)^{\beta-1} \mathbf{1}_{(\mu_k^*, \mu^* + \epsilon]}(\mu), \end{aligned}$$

where $f_k(\mu)$ is the probability density function of arm $k \in K$, $\mathbf{1}_\Theta$ stand for the indicator function of the set Θ , $\gamma_k^1 = 1 - A\epsilon_0^\beta$, $\gamma_k^2 = 1 - \frac{A(\mu^* - \mu_k^* + \epsilon)^\beta}{1 - A\epsilon_0^\beta}$ and γ_k^3 is chosen such that $\int_0^1 f_k^{H_k}(\mu) d\mu = 1$.

Note that since for every $x_1, x_2 \geq 0$ it follows that $x_1^\beta + x_2^\beta \geq (x_1 + x_2)^\beta$ for $\beta \leq 1$, Assumption 1 holds for hypotheses $\{H_1, \dots, H_{|K|}\}$.

To further bound γ_k^2 and γ_k^3 , note that since $\epsilon_0 \leq (4A)^{-1/\beta}$,

$$1 - 2A(\mu^* - \mu_k^* + \epsilon)^\beta \leq \gamma_k^2 \leq 1.$$

Let P_k stands for the mass of an atom in the probability function of arm $k \in K$ at the point $\bar{\mu}_k$ (if there is one), then we note that

$$1 = \int_{-\infty}^{\infty} f_k^{H_k}(\mu) d\mu = \gamma_k^2 (F_k(\bar{\mu}_k) - P_k) + \gamma_k^3 P_k + 1 - F_k(\bar{\mu}_k) + A(\mu^* - \mu_k^* + \epsilon)^\beta \triangleq \Phi(\gamma_k^2, \gamma_k^3),$$

but, since $F_k(\bar{\mu}_k) \geq 1 - A\epsilon_0^\beta$, for $\gamma_k^2 = \gamma_k^3$ it follows that $\Phi(\gamma_k^2, \gamma_k^3) \leq 1$. So, since $\Phi(\gamma_k^2, \gamma_k^3)$ increases in γ_k^3 it is obtained that $\gamma_k^3 \geq \gamma_k^2$. Finally, it follows that in the case of $\epsilon_0 \leq (\mu^* - \mu_k^* + \epsilon)$,

$$\gamma_k \leq \gamma_k^1,$$

and in the case of $\epsilon_0 > (\mu^* - \mu_k^* + \epsilon)$,

$$\gamma_k \leq \min(\gamma_k^2, \gamma_k^3),$$

where

$$\gamma_k \triangleq 1 - 2A(\min(\epsilon_0, \mu^* - \mu_k^* + \epsilon))^\beta$$

If hypothesis H_k ($k \neq 0$) is true, then $\mu_k^* \geq \mu_l^* + \epsilon$ for all $l \neq k$, hence the algorithm should provide a reward from arm k with probability larger than $1 - \delta$. We use E_k^H and P_k^H to denote the expectation and probability, respectively, under the algorithm being considered and hypothesis H_k . Further, for every $k \in K$ let

$$t_k = \frac{1}{4(1 - \gamma_k)} \ln \left(\frac{3}{16\delta} \right),$$

and let T_k stands for the number of samples from arm k .

Suppose now that our algorithm is (ϵ, δ) -correct under H_0 , and that $E_0^H[T_k] \leq t_k$ for some $k \in K$. We will show that this algorithm cannot be (ϵ, δ) -correct under hypothesis H_k . Therefore, an (ϵ, δ) -correct algorithm must have $E_0^H[T_k] > t_k$ for all $k \in K$.

Define the following events:

- $A_k = \{T_k \leq 4t_k\}$. It easily follows from $4t_k(1 - P_0^H(A_k)) \leq E_0^H[T_k]$ that if $E_0^H[T_k] \leq t_k$, then $P_0^H(A_k) \geq \frac{3}{4}$.

- Let B_k stand for the event under which the chosen arm at termination is k , and B_k^C for its complement. Since $P_0^H(B_{k'}) > \frac{1}{2}$ can hold for one arm at most, it follows that $P_0^H(B_k^C) > \frac{1}{2}$ for every $k \in K \setminus \{k'\}$ for some k' .
- Let C_k to be the event under which all the samples obtained from arm k are on the interval $(-\infty, \mu_k^*]$. Clearly, $P_0^H(C_k) = 1$.

Define now the intersection event $S_k = A_k \cap B_k^C \cap C_k$. We have just shown that for every $k \in K \setminus \{k'\}$ it holds that $P_0^H(A_k) \geq \frac{3}{4}$, $P_0^H(B_k^C) > \frac{1}{2}$ and $P_0^H(C_k) = 1$, from which it follows that $P_0^H(S_k) > \frac{1}{4}$ for $k \neq k'$. Further, observe that for every history h_N of N samples for which the event C_k holds, it holds that $\frac{dP_k^H}{dP_0^H}(h_N) \geq (\gamma_k)^N$. We therefore obtain the following inequalities,

$$\begin{aligned} P_k^H(B_k^C) &\geq P_k^H(S_k) = E_0^H \left[\frac{dP_k^H}{dP_0^H} I(S_k) \right] \geq \gamma_k^{-4t_k} P_0^H(I(S_k)) \\ &> \frac{1}{4} \gamma_k^{-4t_k} \geq \frac{1}{4} e^{-\ln \frac{3}{16\delta}} \geq \delta, \end{aligned}$$

where in the last inequality we used the facts that $(1 - \epsilon)^{\frac{1}{\epsilon}} \geq e^{-1}$.

We found that if an algorithm is (ϵ, δ) -correct under hypothesis H_0 and $E_0[T_k] \leq t_k$ for some $k \neq k'$, then, under hypothesis H_k this algorithm returns a sample that is smaller by at least ϵ than the maximal possible reward with probability of δ or more, hence the algorithm is not (ϵ, δ) -correct. Therefore, any (ϵ, δ) -correct algorithm must satisfy $E_0[T_k] > t_k$ for all of arms except possibly for one (namely, for the one k' for which $P_0(B_{k'}^C) \leq \frac{1}{2}$). In addition $t_{k^*} \geq t_{k'}$, where k^* is the optimal arm (namely, $\mu_{k^*}^* = \mu^*$). Hence the lower bound is obtained. \square

8 Appendix B

Proof (Theorem 3). First, we define the following hypotheses:

$$H_0 : f^{H_0}(\mu) = f(\mu),$$

and

$$H_1 : f^{H_1}(\mu) = \gamma f(\mu) + \frac{A}{|K|} \beta (\mu^* + \epsilon - \mu)^{\beta-1} \mathbf{1}_{(\mu^*, \mu^* + \epsilon]}(\mu),$$

where, as in the proof of Theorem 1, $f(\mu)$ is the probability density function of the unified arm, $\mathbf{1}_A$ stand for the indicator function of the set A , and γ is chosen such that $\int_0^1 f^{H_1}(\mu) d\mu = 1$.

Note that since for every $x_1, x_2 \geq 0$ it follows that $x_1^\beta + x_2^\beta \geq (x_1 + x_2)^\beta$ for $\beta \leq 1$, Assumption 1 holds for hypothesis H_1 .

To further bound γ , note that

$$1 = \int_{-\infty}^{\infty} f^{H_1}(\mu) d\mu = \gamma + \frac{A\epsilon^\beta}{|K|},$$

Therefore,

$$\gamma = 1 - \frac{A\epsilon^\beta}{|K|}.$$

If hypothesis H_1 is true, the algorithm should provide a reward greater than μ^* . We use E_l and P_l (where $l \in \{0, 1\}$) to denote the expectation and probability respectively, under the algorithm being considered and under hypothesis H_l . Now, let

$$t = \frac{1}{4(1-\gamma)} \ln \left(\frac{3}{5\delta} \right),$$

and recall that T stands for the total number of samples from the arm.

Now, we assume we run an algorithm which is (ϵ, δ) -correct under H_0 and that $E_0[T] \leq t$ for this algorithm. We will show that this algorithm cannot be (ϵ, δ) -correct under hypothesis H_1 . Therefore, an (ϵ, δ) -correct algorithm must have $E_0[T] > t$.

Define the following events:

- $A = \{T \leq 4t\}$. By the same consideration as in the proof of Theorem 1 (for the events $\{A_k\}_{k \in K}$), it follows that if $E_0[T] \leq t$, then $P_0(A) \geq \frac{3}{4}$.
- Let B stand for the event under which the chosen sample is smaller or equal to μ^* , and B^C for its complementary. Clearly, $P_0(B) = 1$.
- We define the event C to be the event under which all the samples obtained from the unified arm are on the interval $[-\infty, \mu^*]$. Clearly, $P_0(C) = 1$.

Define now the intersection event $S = A \cap B^C \cap C$. We have shown that $P_0(A) \geq \frac{3}{4}$, $P_0(B) = 1$ and $P_0(C) = 1$, from which it is obtained that $P_0(S) \geq \frac{3}{4}$. In addition, since for every history h_N of N samples, for which the event C holds, it is obtained that $\frac{dP_1}{dP_0}(h_N) \geq \gamma^N$, we have the following,

$$\begin{aligned} P_1(B) &\geq P_1(S) = E_0 \left[\frac{dP_1}{dP_0} I(S) \right] \geq \gamma^{-4t} P_0(I(S)) \\ &\geq \frac{3}{4} \gamma^{-4t} \geq \frac{3}{4} e^{-\ln \frac{3}{5\delta}} \geq \delta, \end{aligned}$$

where in the last inequality we used the facts that $(1 - \epsilon)^{\frac{1}{\epsilon}} \geq e^{-1}$.

We found that if an algorithm is (ϵ, δ) -correct under hypothesis H_0 and $E_0[T] \leq t$, then, under hypothesis H_1 this algorithm returns a sample that is smaller by at least ϵ than the maximal possible reward with a probability of δ or more, hence the algorithm is not (ϵ, δ) -correct. Therefore, any (ϵ, δ) -correct algorithm, must satisfy $E_0[T] > t$. Hence the lower bound is obtained. \square

9 Appendix C

Proof (Theorem 4). Since sampling from the unified arm consists of choosing one arm out of the $|K|$ arms (with equal probability), and then, sampling from this arm, it follows that, $F(\mu^* - \epsilon) \leq \left(1 - \frac{A\epsilon^\beta}{|K|}\right)$. Also, we note that $(1 - \epsilon)^{\frac{1}{\epsilon}} \leq e^{-1}$ for every $\epsilon \in (0, 1]$. Therefore, for $N = \lceil \frac{-\ln(\delta)|K|}{A\epsilon^\beta} \rceil + 1$,

$$P(V_N^1 < \mu^* - \epsilon) = (F(\mu^* - \epsilon))^N \leq \left(1 - \frac{A\epsilon^\beta}{|K|}\right)^N < \delta, \quad (13)$$

where V_N^1 is the largest reward observed among the first N samples. Hence, the algorithm is (ϵ, δ) -correct. The bound on the sample complexity is immediate from the definition of the algorithm. \square